

Large Scale Signal Detection: A Unified Perspective

Subhadeep Mukhopadhyay

Department of Statistical Science, Temple University

Philadelphia, Pennsylvania, 19122, U.S.A.

deep@temple.edu

Dedicated to John Tukey, on the occasion of his 100th birthday.

ABSTRACT

There is an overwhelmingly large literature and algorithms already available on ‘large scale inference problems’ based on different modeling techniques and cultures. Our primary goal in this paper is *not to add one more new methodology* to the existing toolbox but instead (a) to clarify the mystery how these different simultaneous inference methods are *connected*, (b) to provide an alternative more intuitive derivation of the formulas that leads to *simpler* expressions in order (c) to develop a *unified* algorithm for practitioners. A detailed discussion on representation, estimation, inference, and model selection is given. Applications to a variety of real and simulated datasets show promise. We end with several future research directions.

Keywords: Mixed-sample problem, False discovery rate; Comparison density; RKHS; Smooth p-value; Skew-beta density decomposition; Pre-flattening smoothing; Tail modeling.

1 Introduction

1.1 Mixed Sample Inference

Consider n i.i.d. observation Z_1, \dots, Z_N from the continuous distribution $F = \pi_0 F_0 + (1 - \pi_0)G$, $0 < \pi_0 \leq 1$, where F_0 is specified distribution (known) and G is unknown mixing distribution. As the given random sample may contain observations from both F_0 and G , we call this framework a ‘mixed-sample problem.’ This can be considered as an intermediate statistical inference problem between two extremes: one and two sample inference problems. Typical goals of mixed data inference problems include: (1) identifying the Z_i ’s coming from the distribution G (also known as signal detection), and (2) estimating the proportion π_0 .

Motivation. A multiple hypothesis testing problem can be thought of as a mixed-sample problem by considering Z_i 's to be the test statistic corresponding to a null hypothesis H_{0i} ($i = 1, \dots, N$), where the goal is simultaneous inference. Our main motivation is to develop a comprehensive modeling framework for big- N signal detection (or large-scale mixed-sample) problems in a way that is *analogous* to the one-sample and two-sample approaches so that we can cover the whole statistical inference spectrum (one sample \rightarrow mixed sample \rightarrow two sample) using one general concept and tool – which will simplify theory, practice and teaching.

1.2 Goals and Organization

The topic of multiple hypothesis testing, which began with the pioneering work of Tukey on “Comparing individual means in the analysis of variance” published in 1949 *Biometrics*, gained new momentum in the 21st century with modern high-throughput data acquisition. Since then an enormous amount of research has been conducted on this topic. Following is a summary of our research contributions:

(1) We seek to *condense and systemize* the vast literature on large-scale simultaneous inference methods based on a few fundamental concepts and notations introduced in Section 2. The author believes that this could ultimately help applied data scientists to gain better insight for selecting the proper signal detection tool from the existing large inventory by simplifying the practice. Section 2.2 describes a *unified representation theory and an alternative rationale* of three main classes of large-scale inference procedures: Benjamini Hochberg’s false discovery rate (FDR) approach (Benjamini and Hochberg, 1995), Efron’s empirical Bayes local FDR proposal (Efron et al., 2001), and Donoho Jin’s Higher Criticism thresholding (Donoho and Jin, 2004).

(2) Section 2.2 further shows how this new theoretical framework for large scale simultaneous inference problems can *integrate and reconcile* frequentist and Bayesian cultures, *bypassing* the old philosophical and ideological differences (Benjamini, 2008) to develop unified efficient algorithms using single concept and notation.

(3) The theory presented in Section 2 motivated us to introduce a *new class of nonparametric signal detection algorithm* by converting the large-scale inference problems into a function estimation problem based on the ideas of comparison density and distribution (CD); see

Web Appendix A for more discussion on the historical significance of the current research. Section 3.1 describes the statistical modeling challenges. A comprehensive CD-based non-parametric modeling solution for ‘signal-hunting’ procedure is described in Section 3.2-3.5, which consists of three main components: (A) We discover that comparison density takes a ‘*universal stylized shape*’ for almost all large-scale inference problems - ‘U’ shaped density over the compact unit interval. To tackle this nonstandard density we propose a specially designed Skew-Beta decomposition-based nonparametric estimation algorithm; (B) To model empirical null we propose a robust biweight M-estimator, which is easy to compute; (C) Finally, we describe the Minimum Deviance Criteria (MDC) algorithm for estimating true null proportion. Section 4 presents numerical results, and discussions are given in Section 5.

2 Theory

2.1 Background on Functional Statistical Inference

Let Z_1, Z_2, \dots, Z_N be a mixed random sample, with the majority of the observations coming from continuous null distribution F_0 , and a small proportion from unknown contaminating distribution G : $F = \pi_0 F_0 + (1 - \pi_0)G$, $0 < \pi_0 \leq 1$, where Z_i ’s are the corresponding test statistic for the hypothesis testing problem H_{0i} and the goal is simultaneous inference. Here we will provide an intuitive introduction to the functional statistical approach for large-scale simultaneous inference problems whose goal is to detect false null hypotheses.

The cornerstone of our approach is based on the concept of comparison density (Parzen, 1983). For continuous F and G , define comparison distribution function $D(u; G, F) := F(G^{-1}(u))$ and the corresponding comparison density by:

$$d(u; G, F) = \frac{f(G^{-1}(u))}{g(G^{-1}(u))}, \quad 0 < u < 1. \quad (2.1)$$

Under $H_0 : F = F_0$, i.e., when all the observations actually come from F_0 , the theory of weak convergence of comparison distribution process (Parzen, 1999, Thas, 2010) tells us $\sqrt{N}(\tilde{D}(u) - u) \xrightarrow{d} \mathbb{B}(u)$ as $N \rightarrow \infty$, where $\mathbb{B}(u)$, $0 < u < 1$ is the Brownian bridge process, and $\tilde{D}(u) = \tilde{F}(F_0^{-1}(u))$, where $\tilde{F}(z; Z) = N^{-1} \sum_{i=1}^N \mathbb{I}(Z_i \leq z)$. In mixed-sample problems, our goal is to identify the signals or the Z_i ’s that are coming from G . If all the null-hypothesis $H_{01}, H_{02}, \dots, H_{0N}$ are true (i.e., under $H_0 : F = F_0$), we would expect

$\tilde{D}(u) := D(u; F_0, \tilde{F}) \approx u$. Thus, intuitively, the collection of u 's, for which the distance between $\tilde{D}(u)$ and u are large, should be suspected as potential signals. In other words,

$$\{u : \tilde{D}(u)/u > \gamma\},$$

for a suitably selected threshold might contain the signals or the rejected null hypotheses. Now instead of comparison distribution function $D(u)$, in the same spirit, we can develop a similar strategy for detecting false null hypotheses based on comparison density $d(u)$. So the performance of simultaneous inference critically depends on the fundamental comparison density (or equivalently distribution) function.

From Multiple Testing to Nonparametric Function Estimation. The crux of this section is this: the notion of comparison distribution and density allows us to transform the simultaneous hypothesis testing problem into a nonparametric function approximation problem - a functional statistical approach to large scale inference. In the next section, we will make this intuitive mixed-sample signal detection approach formal and rigorous.

2.2 Unified Representation and Theory of Threshold Selection

The previous section intuitively argued how to build a CD-based simultaneous inference algorithm. The purpose of this section is threefold: (1) to derive proper thresholds for CD-based simultaneous inference procedures that can lead to desirable methods with provable guarantees; (2) to show how the Frequentist and Bayesian large-scale inference algorithms can be connected using the theory of reproducing kernel Hilbert space (RKHS) of the limiting Brownian bridge process of the comparison distribution; (3) to simplify the currently existing vast literature on large-scale multiple testing by providing unified representation based on a single concept and notation – comparison density/distribution, to our knowledge largely an unexplored area of research.

Let's start with the question of how to select the proper threshold. The answer depends on the appropriate error rate that we seek to control in multiple hypothesis testing. Based on particular set up and purpose of any specific inference problems Benjamini (2010) provided a list of 14 such error rates.

One of the important error-rate criteria (among many possibilities) in large- N multiple testing is to control the false discovery rate (FDR) – the expected proportion of Type I errors among the rejected hypotheses, suggested by Benjamini and Hochberg (1995). The

following theorem describes how to select the threshold for comparison distribution-based method that can guarantee $\text{FDR} \leq \alpha$ (for some preselected significance level α) under any arbitrary π_0 and G . A sketch of the proof is provided in Web Appendix B.

Theorem 1. *Consider testing N independent null hypothesis $H_{(01)}, \dots, H_{(0N)}$ based on corresponding ordered p -values $u_{(1)}, \dots, u_{(N)}$. Define the index set*

$$\mathcal{R} = \left\{ i \leq k : k = \operatorname{argmax}_i \frac{\tilde{D}(u_{(i)})}{u_{(i)}} \geq \frac{\eta}{\alpha} \right\}. \quad (2.2)$$

Then the procedure that rejects $H_{(0i)}, i \in \mathcal{R}$ controls FDR at the level α , regardless of the distribution of the test statistic corresponds to false null hypothesis.

To see how CD-based algorithm (2.2) essentially equivalent to the (adaptive) Benjamini Hochberg's (BH) FDR controlling procedure, note that $\tilde{D}(u_i; F_0, \tilde{F}) = \tilde{F}F_0^{-1}(u_{(i)}) = i/N$, which is equivalent to saying that reject $H_{(0i)}$ $i = 1, \dots, k$, where

$$k = \max \left\{ 1 \leq i \leq N : u_{(i)} \leq \frac{i}{N} \frac{\alpha}{\pi_0} \right\}.$$

RKHS-based dual representation. Here we will give an alternative CD-based FDR controlling procedure motivated by the isometric isomorphism (one-one inner product preserving map) relation between the Hilbert space spanned by the Brownian bridge process (which is the limit process of $\sqrt{N}(\tilde{D}(u) - u)$ and the corresponding RKHS associated with the covariance kernel $K_{\mathbb{B}}(u, v)$. The following theorem provides the complete characterization of the RKHS.

Theorem 2. *The reproducing kernel Hilbert space H_K associated with the Brownian bridge covariance kernel $K_{\mathbb{B}}(u, v) = \min(u, v) - uv$, $0 < u, v < 1$ consists of L^2 differentiable functions with the inner produce $\langle \phi, \psi \rangle = \int_0^1 \phi'(u)\psi'(u) du$ that satisfies $\phi(0) = \phi(1) = 0$.*

It is sufficient to prove the reproducing property relative to the inner product in the following sense, for all $\phi \in H_K$

$$\langle K_{\mathbb{B}}(u, \cdot), \phi \rangle = \phi(u), \text{ for any } u \in [0, 1]. \quad (2.3)$$

For details of the proof see Web Appendix C.

Moving toward Empirical Bayes Formulation. As the RKHS associated with the Brownian bridge covariance kernel equipped with the norm squared $\|h\|^2 = \int_0^1 |h'(u)|^2 du$, this tells us that the other bonafide signal detection technique can equivalently be constructed by

looking at the distance between $D'(u) = d(u)$ and 1. That is, the collection of u 's for which $d(u; F_0, F)$ *substantially* deviates from $\text{Uniform}[0, 1]$, are precisely the non-null candidates that we are searching. A comparison density $d(u; F_0, \tilde{F})$ based signal characterization (instead of comparison distribution function $D(u; F_0, \tilde{F})$) is given in the following result, which turns out to be an alternative way of expressing Efron's empirical Bayes local fdr (Efron et al., 2001) formula.

Theorem 3. *Local false discovery rate can alternatively be represented using the comparison density*

$$\text{fdr}(z) := \Pr\{\text{null} \mid Z = z\} = \frac{\pi_0}{d(F_0(z); F_0, F)}. \quad (2.4)$$

The local fdr is defined as the conditional probability of a case being null or noise given $Z = z$,

$$\text{fdr}(z) = \Pr(\text{null} \mid Z = z) = \pi_0 \frac{f_0(z)}{f(z)} = \pi_0 / d(F_0(z); F_0, F). \quad (2.5)$$

It is known that the local FDR is more conservative than the BH procedure. To address this issue, Efron (2007) recommended the threshold $\{u : d(u; F_0, \tilde{F}) > \tau_{\text{BH}}/2\}$ (under the Lehmann alternatives), where $\tau_{\text{BH}} = \pi_0/\alpha$, which he argued controls size, or Type I errors at the desired level. The goal of local FDR algorithm is to estimate the conditional probability $\Pr(\text{null} \mid Z = z)$ from the data. The traditional procedure estimates *separately* $\hat{\pi}_0$, $\hat{f}_0(t)$ and $\hat{f}(t)$, while our comparison density-based approach *directly* estimates the probability; see Web Appendix D for more discussion on this point.

Relation with Higher Criticism. Another related method for detecting rare/weak signals in large-scale experiments is ‘‘Higher Criticism’’ (HC) thresholding (which is not primarily designed to protect against false discoveries) introduced by Donoho and Jin (2004) generalizing Tukey’s proposal (Tukey, 1989). Following is the equivalent representation of HC-procedure using our notation. Reject $H_{(0i)}$ $i = 1, \dots, k$ where

$$k = \left\{ 1 \leq i \leq \alpha_0 N : \arg\max_i \frac{\tilde{D}(u_{(i)}) - u_{(i)}}{\sqrt{u_{(i)}(1 - u_{(i)})}} \right\}, \quad (2.6)$$

where $\alpha_0 = .1$ is a common choice. We can study the asymptotic properties of HC procedure using the limiting unit variance Gaussian process of $\sqrt{w(u)}(\tilde{D}(u) - u)$, $w(u) = 1/u(1 - u)$ with the corresponding covariance kernel $K(u, v) = \frac{\min(u, v) - uv}{\sqrt{u(1 - u)v(1 - v)}}$; See Donoho and Jin (2004) for further details.

A New Class of Nonparametric Large-scale Inference Procedure. The representation results proved in this section allow us to develop a new class of CD-based approaches by specifying different models and estimation strategies for \hat{d} (or equivalently $\hat{D}(u) = \int_0^u \hat{d}(v) dv$). For example, the *raw-empirical* comparison distribution function \tilde{D} generates the BH procedure (2.2) or the HC procedure (2.6), which can be upgraded to a more enhanced technique by *plugging in smooth nonparametric estimate* $\hat{D}(u) = \int_0^u \hat{d}(v) dv$.

Nonparametric Function Approximation. We argue comparison density is an indispensable tool that provides the required theoretical foundation to develop a single general algorithm for modern large-scale signal detection problems. The CD-based approach unifies the frequentist BH procedure and empirical Bayes local false discovery concepts by converting the simultaneous inference problem into a nonparametric function estimation problem.

Any comprehensive multiple testing algorithm should have three main components: (1) π_0 (null-proportion), (2) F_0 (empirical null distribution), and (3) $d(u; F_0, F)$ (comparison density). The efficiency of the signal detection procedure directly depends on the quality of estimation of these quantities. In the next section, we will describe a nonparametric estimation algorithm (CDfdr) in a step-by-step manner (along with real data illustration) to achieve this goal.

The following section starts with an interesting observation that virtually all large-scale inference problems generate comparison density with a very *special universal shape* (see Web Appendix K for more discussion on the tail-behaviour of mixed-sample comparison density). It turns out that the traditional nonparametric density estimation methods fail to satisfactorily capture this typical shape of the comparison density over a compact $[0, 1]$ support. For that purpose, we develop a *specially designed nonparametric algorithm* (tailored for this particular characteristic shape) for parsimonious comparison density estimation \hat{d} , which will be used for detecting signals or false null hypotheses and to estimate the null proportion π_0 .

3 Estimation

3.1 Two Modeling Challenges

Fig 1(A) shows the two-sample t-test p-values of $N = 7129$ gene expressions of Golub cancer data (Golub et al., 1999) while comparing $n_1 = 27$ acute lymphoblastic leukemia (ALL) and $n_2 = 11$ acute myeloid leukemia (AML) tumor samples to identify differentially expressed genes (data available in R package `golubEsets`). The comparison density, which is the distribution of p-values from this large-scale study indicates two nonparametric modeling challenges: First, we need density estimators for compact support $[0, 1]$. Second, the sharp narrow peak near the boundary 0 necessitates modeling of the highly dynamic tail.

Developing nonparametric density estimation for compact support $[0, 1]$ is known to be a challenging problem due to the “boundary effect.” Recently, there has been a great deal of interest in adapting kernel density estimator for unit interval (see Geenens (2014) and Wen and Wu (2014)) by transforming the variable of interest into another one whose density estimation should be free from boundary problems, and then finally transforming that estimate back into the initial scale. However, it has been recognized that these methods are computationally not efficient. Other approaches, like regression-based density estimators via smoothing splines or local polynomials, are known to have a larger variance near the boundaries (Thas, 2010).

For large-scale signal detection problems, besides compact support, the other more challenging modeling aspect is to tackle the highly dynamic tail near the boundaries 0 and 1, where all previously mentioned density estimators perform poorly. We will propose a specially tailored *new genre of density estimator* to tackle this typical non-standard “U” shaped density on the unit interval.

We do note, however, that, it is not difficult to propose new density estimation techniques to fit the data, such as Fig 1(A), but it is less trivial to come up with a sparse parametric model that fits the data well and is easy to interpret. The most stunning fact about our approach (that we will elaborate on the next section) is that it require only *three parameters* to accurately model the golub p-values, including the tail region!

To model this typical shape of comparison density (that arises in the large scale signal

detection problems), we prefer estimators that are simultaneously accurate, computationally simple and parsimonious. Our proposal consists of two main steps: (1) converting “spiky” p-values to “smooth” p-values via the preflattening technique as described in Section 3.2 and (2) estimating “smooth” p-values by expanding preflattened $d(u)$ (nonparametric L^2 or orthogonal series density estimator) or $\log d(u)$ (maximum entropy exponential density estimator) as a linear combination of orthonormal shifted Legendre polynomials (whose support is $[0, 1]$), described in Section 3.3. *The novelty of our approach lies in its unique ability to “decouple” the density estimation problem into two separate modeling problems: the tail part and the central part of the distribution.*

[Figure 1 about here.]

3.2 Skew-Beta Density Estimator: Tail Modeling

We propose a new genre of nonparametric density estimation technique, which starts with a parametric model $g(x)$ and permits the following *universal decomposition*:

$$f(x) = g(x) \times d[G(x); G, F]. \quad (3.1)$$

Verify comparison density (2.1) evaluated at $G(x)$ has the form $d(G(x); G, F) = f(x)/g(x)$. We call this new class as Skew-G density model for $f(x)$.

To efficiently capture the typical shape of the p-values shown in Fig 1(A), we select $g(x)$ on $[0, 1]$ in such a way that it can tackle the rapidly changing tail. Hence, beta distribution is a good candidate for $g(x)$, which can act as pre-flattening function, as shown in Fig 1(B,C). The resulting skew-beta comparison density estimator has the following decomposition:

$$d(u; F_0, F) = f_B(u; \alpha, \beta) \times d[F_B(u; \alpha, \beta); F_B, F], \quad \text{for } 0 < u < 1, \quad (3.2)$$

where beta density and distribution is denoted by f_B and F_B . We define the quantity $F_B(u; \alpha, \beta)$ as the “smooth” p-value, shown in Fig 1(C), whose density is $d[F_B(u; \alpha, \beta); F_B, F]$, $0 < u < 1$. Use simple maximum-likelihood estimate $\hat{\alpha}$ and $\hat{\beta}$ to get the fitted beta; see Web Appendix E for more discussion and simple one line R implementation. Note that if none of the genes were differentially expressed for Golub data then we would expect the distribution of the p-values to be $\text{Beta}(\alpha = 1, \beta = 1)$, i.e., uniform. This simple observation can lead to a quick diagnostic for testing the presence of signals. Web Appendix F describes the procedure.

3.3 Estimating density of “Smooth” P-values

In this Section, we develop a nonparametric density estimator for smooth (Beta-transformed) p-values $d(v; F_B, F)$ on the unit interval, where $v = F_B(u; \alpha, \beta)$. Our proposal is based on L^2 or maximum entropy exponential comparison density estimator by expanding it in an orthonormal basis. We select shifted orthonormal Legendre Polynomial (LP) $\text{Leg}_j(v)$, $j = 1, 2, \dots$ as basis, which form a complete orthonormal basis for $L^2[0, 1]$ Hilbert space of square-integrable functions.

$$L^2 \text{ orthogonal series estimator: } d(v; F_B, F) - 1 = \sum_j \text{LP}[j; F_B, F] \text{Leg}_j(v) \quad (3.3)$$

$$\text{Maximum entropy estimator: } \log d(v; F_B, F) = \sum_j \theta_j \text{Leg}_j(v) - K(\boldsymbol{\theta}) \quad (3.4)$$

Due to simplicity and computational ease (which will be clear soon) we use L^2 estimate, which has nice asymptotic theoretical properties (Anderson and de Figueiredo, 1980). The large- N paradigm (for golub data $N = 7129$) makes all of these asymptotic analyses very much relevant and directly applicable. However, the reader is free to use the nonparametric maximum entropy exponential estimate (3.4).

Computation of LP-Fourier Coefficients. The L^2 orthogonal expansion coefficients satisfies

$$\text{LP}[j; F_B, F] = \int_0^1 d(v; F_B, F) \times \text{Leg}_j(v) dv. \quad (3.5)$$

Substitute v by $F_B(u; \alpha, \beta)$ and note

$$d[F_B(u; \alpha, \beta); F_B, F] = \frac{d(u; F_0, F)}{f_B(u; \alpha, \beta)}.$$

This lead to the following representation:

$$\text{LP}[j; F_B, F] = \int_0^1 d(u; F_0, F) \text{Leg}_j \{F_B(u; \alpha, \beta)\} du. \quad (3.6)$$

Eq. (3.6) implies that the LP-coefficients can be represented as the mean of Leg_j *evaluated at* the $F_B(u; \alpha, \beta)$ (these are Beta-transformed smooth pvalues). This representation result is summarized in the following theorem.

Theorem 4. *The Fourier-LP coefficients admit the following representation*

$$\text{LP}[j; F_B, F] = \int_0^1 \text{Leg}_j(v) dD(v; F_B, F) = \mathbb{E}[\text{Leg}_j\{F_B(U; \alpha, \beta)\}].$$

As a practical consequence, this result allows fast computation of the LP expansion coefficients as mean of the Leg_j score functions *evaluated at the transformed smooth pvalues*:

$$\text{LP}[j; F_B, \tilde{F}] \leftarrow N^{-1} \sum_{i=1}^N \text{Leg}_j [F_B(u_i; \alpha, \beta)]. \quad (3.7)$$

Adaptive Estimation. Estimate the smooth nonparametric model by selecting the ‘significantly large’ LP-coefficients in a data-driven way. We will use the Ledwina (1994) scheme using Schwarz selection criterion, which is known to be consistent. To identify the important coefficients using Ledwina’s data-driven model selection criterion, we first rank the squared LP-coefficients. Then we take the penalized cumulative sum of k coefficients using information criterion $N^{-1} \log(N)k$ and choose the k for which this is the maximum. Fig 1(D) shows the resulting smooth estimate for the golub data.

3.4 Estimating Proportion of True Null Hypothesis

In Section 2.2 we have seen the thresholds for CD-based false discovery methods depend on the parameter π_0 , the true proportion of noise or null hypothesis. The following is our proposed data-analytic estimation algorithm.

Algorithm 1 [π_0 estimation by Minimum Deviance Criteria (MDC)]

Step 1. Define $\mathcal{U}_\lambda = \{u_i : \hat{d}(u; F_0, F) < \lambda\}$, where \hat{d} is the estimated beta-preflattened nonparametric comparison density; $|\mathcal{U}_\lambda| = N_\lambda$. For each fixed λ perform the following steps.

- (a) Compute $\widetilde{\text{LP}}_\lambda[j] \leftarrow N_\lambda^{-1} \sum_{i=1}^{N_\lambda} \text{Leg}_j(u_i)$, which is the score coefficient for the following L_2 comparison density $\tilde{d}_\lambda(u) = 1 + \sum_{j=1}^M \widetilde{\text{LP}}_\lambda[j] \text{Leg}_j(u)$, based on \mathcal{U}_λ .
- (b) Calculate the deviance statistic $I_\lambda \leftarrow \sum_{j=1}^M |\widetilde{\text{LP}}_\lambda[j]|^2$.

Step 2. Display the deviance path (λ, I_λ) on a fine grid of $\lambda \in [1, \gamma]$ and set $\lambda^* \leftarrow \arg \min_\lambda I_\lambda$.

Step 3. Output $\hat{\pi}_0 = N_{\lambda^*}/N$.

The *rationale* behind the algorithm comes from the simple fact that under H_0 , when all the cases are null, the underlying comparison density $d(u; F_0, F)$ should not deviate much from $\text{Uniform}[0, 1]$ as $D(u; F_0, F) = u$. Parseval’s theorem dictates the following equality for L^2

comparison density:

$$\int_0^1 [d(u; F_0, F) - 1]^2 du = \sum_j |\text{LP}[j; F_0, F]|^2. \quad (3.8)$$

In the light of Eq. (3.8) the statistic I_λ can be interpreted as the deviation of $\tilde{d}_\lambda(u)$ from uniformity. Fig 1 of Web supplement illustrates this idea for Prostate cancer dataset (described in Section 4), where the shape of the estimated comparison density clearly indicates the presence of signals in the two tails. The deviance path for the rejection region of interest $1 \leq \lambda \leq \gamma = 3.5$ is shown in the right panel for $M = 10$, which gives $\arg \min_\lambda I_\lambda = 1.98$ and $\hat{\pi}_0 = 0.971$. At the point $\lambda^* = 1.98$ the deviance statistics takes the minimum value, which implies that the set $\mathcal{U}_{\lambda=1.98}$ is most likely consists of the null cases.

Connection with other approaches. Instead of working with comparison density, there are related methods based on comparison distribution function to identify p-values that differ significantly from the null Uniform[0, 1]. Our method could be viewed as the formal density analogue of the graphical method proposed in Schweder and Spjøtvoll (1982). Another similar method proposed by Storey (2002) estimates $\hat{\pi}_0(\lambda) = (1 - \tilde{D}(\lambda))/(1 - \lambda)$. They proposed a computationally intensive method to select the tuning parameter λ . It consists of bootstrapping p-values for each λ in the range $\{0, .05, \dots, .95\}$ and selecting the one which minimizes mean-squared error (MSE).

3.5 Estimating Empirical Null

Until this point we have assumed that the continuous null distribution F_0 is completely specified. As noted by Efron (2010), “This is a less-than-usual occurrence” for real datasets. Efron (2004) proposed a method for estimating the unknown location and scale parameters under the zero assumption (i.e., the non-null or the mixing distribution (signal) $g(z) = 0$ in a certain interval around the $z = 0$, which is known to be prone to bias) for $F_0 = \Phi[(x - \mu_0)/\sigma_0]$. Muralidharan (2010) estimates by putting a Dirichlet(β) prior penalty on the π , where β is an extra tuning parameter. Here we propose a simple and fast solution that works quite nicely.

In the event all the observations Z_1, \dots, Z_N are coming from F_0 , we have $F(z) = F_0[(z - \mu_0)/\sigma_0]$, where F_0 is the theoretical location-scale null distribution. We could have easily estimated the unknown μ_0 and σ_0 just by fitting a simple linear regression model in the

QQ plot. But instead, we have few Z_i 's, which are coming from G (π_0 -contamination neighborhood of F_0). If we knew which of these were unusual observations we could have removed them before fitting the linear regression. But as this is not the case, the natural step is to fit *robust* linear regression instead of linear least square, which is resistant to the data contamination. To get the estimates, solve the M-estimating equation:

$$\begin{aligned}\sum_{i=1}^n \Psi_T \left[\tilde{F}^{-1}(u_i) - \mu_0 - \sigma_0 F_0^{-1}(u_i) \right] &= 0, \\ \sum_{i=1}^n \Psi_T \left[\tilde{F}^{-1}(u_i) - \mu_0 - \sigma_0 F_0^{-1}(u_i) \right] F_0^{-1}(u_i) &= 0,\end{aligned}\tag{3.9}$$

where we choose Ψ_T to be the Tukey's biweight influence function (Beaton and Tukey, 1974) given by

$$\Psi_T(z) = \begin{cases} z[1 - (z/k)^2]^2 & \text{for } |z| \leq k, \\ 0 & \text{for } |z| > k. \end{cases}$$

The value $k = 4.685$ is usually used as a default choice that provides 95% asymptotic efficiency under normality and still offers protection against outliers. A straightforward R-implementation is possible via MASS library function `rlm`; see Web Appendix G for R-code. Table 1 reports the findings of our approach for five large-scale studies discussed in Chapter 6 of Efron (2010). Finally, produce the null-adjusted nonparametric comparison density estimate using $\hat{F}_0 := F_{0;\hat{\mu}_0,\hat{\sigma}_0}$

$$\hat{d}(u; \hat{F}_0, F) = f_B(\hat{F}_0; \hat{\alpha}, \hat{\beta}) \times \hat{d}(F_B(\hat{F}_0; \hat{\alpha}, \hat{\beta}); F_B, F).\tag{3.10}$$

[Table 1 about here.]

4 Examples

4.1 Real Data Application

Prostate cancer data (Singh et al., 2002) consists of 102 patient samples (50 labeled as normal and 52 as prostate tumor samples) and 6033 gene expression measurements. We aim to detect interesting genes that are differentially expressed in the two samples. For this purpose, we compute the two-sample t-test statistic t_i for each gene and convert them into z-scale by $z_i \leftarrow \Phi^{-1}(\mathcal{T}_{100}(t_i))$, where \mathcal{T} denotes the t-distribution function, shown in Fig 2(A) of Web Supplement. Fig 2(D) of Web Supplement shows the final beta-preflattened

smooth estimate of comparison density given by

$$\widehat{d}(u; \Phi, F) = .68 [1 + 0.057 \text{Leg}_6 (F_B(u; \widehat{\alpha} = .81, \widehat{\beta} = .82))] u^{-.19} (1 - u)^{-.18}, \quad 0 < u < 1, \quad (4.1)$$

which along with the Minimum Deviance Criteria gives $\widehat{\pi}_0 = .971$. The representation result (2.5) given in Theorem 3 immediately provides a CD-based estimate of local fdr (which we call **CDfdr**). We compare our result with Locfdr and Mixfdr (Muralidharan, 2010) that estimate the fdr by separately estimating the numerator \widehat{f}_0 and the denominator \widehat{f} (see Web Appendix D). Naturally there are many variants available for these two methods depending on the way they estimate null and marginal densities. We have used the R package **locfdr** and **mixfdr** for implementation purpose. Locfdr estimates pool density f using splines. Methods for estimating null include the following: (a) theoretical ($\mathcal{N}(0, 1)$); (b) maximum likelihood (MLE); (c) central matching (CM); (d) split-normal (SN). Mixfdr implement J group normal mixture model for f . Estimation of empirical null involves putting Dirichlet prior on mixing proportion. We have used the default choice of J and Dirichlet parameter P throughout. All the three empirical null methods, including ours (reported in Table 1) shows that the theoretical null is not drastically different from $\mathcal{N}(0, 1)$. A very slight scale correction was required.

[Figure 2 about here.]

Each the empirical null based estimated fdr functions, including ours, shows strong similarity in the tail-regions. The behavior of the fdr function in the tail-region (compared to the central region) is most crucial, as this determines which p-values will be declared as false or non-null. So, it might be more appropriate to focus on the quality of “tail-modeling” instead of considering the entire curve. We carry out this in the next section where we carefully quantify the estimation accuracy, especially in the tails.

The number of non-null genes identified by different methods, along with the estimates of the proportion of true null is provided in the Web Supplement Table 2. First note that the estimates of π_0 from Locfdr-CM and Locfdr-SN are unrealistic as they exceed 1. Number of non-null genes identified using CDfdr matches with the Locfdr-SN and Mixfdr-Emp method. Furthermore, our proposed minimum deviance estimate $\widehat{\pi}_0$ is very close to the Mixfdr-Th, Mixfdr-Emp and Locfdr-Th methods.

4.2 Simulation Study

In order to further evaluate the accuracy of the **CDfdr** algorithm, we perform two simulated experiments. We are mainly interested in investigating how accurately different methods estimate fdr , especially in the tail-regions based on the mean integrated square error (MISE) criteria. Comparisons will be done with **Locfdr**, **Mixfdr** and **Fdrtool** (Strimmer, 2008). Grenander density estimation is used in **Fdrtool** for estimating the unconditional density f and is implemented in R package **Fdrtool**.

4.2.1 Mixture Normal

We simulate $T_i \sim \mathcal{N}(\mu_i, 1), i = 1, \dots, N = 5000$ out of which 4500 μ_i is set to zero. The remaining 500 is drawn from (once and for all) $\mathcal{N}(\mu, 1)$. We estimate the $\widehat{\text{fdr}}$ for various methods and repeat the whole process 150 times for $\mu = 0.2, 0.5, 1, 2$. Our setup closely follows Storey (2002) and Muralidharan (2010).

Local fdr estimation. The goal is to investigate the efficiency of different estimation methods when we fix the null density at $\mathcal{N}(0, 1)$. Web Supplement Figure 3 depicts the expectation and standard deviation of local fdr for various methods under four different choices of μ . For $\mu = 2$, when the signal and noise are well-separated, all of the methods perform equally well apart from **Fdrtool**, which not only shows high bias but has large variability in the crucial tail region. Under the more difficult scenario of $\mu = .2$, clearly **CDfdr** is the only method that can claim to be unbiased. The variability of **Mixfdr** and **CDfdr** seems similar, though in the extreme (right) tail, **CDfdr** shows more stability. If we look at the $\text{Sd}(\widehat{\text{fdr}})$ curves for **CDfdr** and compare with other competing methods, it appears to be the least variable. Also, **CDfdr** achieves *near unbiasedness irrespective of the underlying signal strength*, which makes it a reliable tool for large-scale inference problems.

Null proportion estimation. It is interesting to examine how all of these methods would perform in estimating the true null proportion under different degrees of signal strength. We implemented our Minimum Deviance Criteria (MDC) algorithm to estimate π_0 . The simulation was done under the same experimental setup. The boxplots of estimates of π_0 under different non-null densities are shown in Web Supplement Figure 4. **Locfdr** shows the largest variability. **Mixfdr** certainly performs best among the competing methods. However, the method that was particularly successful for estimating the true null proportion $\pi_0 = .9$

quickly and accurately is based on the our proposed MDC. This makes the CDfdr algorithm more powerful and efficient as a signal detection tool that protects against false discovery.

4.2.2 Mixture Uniform

We generate p-values from the following model with the parameter choices: $\pi_0 = \{0.9, 0.95, 0.99\}$ and $a = \{0.02, 0.002\}$:

$$\pi_0 \text{Uniform}[0, 1] + (1 - \pi_0) \text{Uniform}[0, a]. \quad (4.2)$$

Here we particularly pay attention to the tails and for that we consider the tail-specific MISE criteria $\mathbf{E} \int_{\mathcal{S}} (\text{fdr}(u) - \hat{\text{fdr}}(u))^2 du$, where \mathcal{S} denotes the collection of u coming from the alternative model $U[0, a]$. The goal is to quantify how precisely the fdr is estimated for the signals (in the tail). Here the parameter a controls the signal strengths and parameter π_0 determines the underlying sparsity levels.

[Figure 3 about here.]

Our simulation design covers the complete spectrum from

dense and weak \rightarrow rare and weak \rightarrow strong and dense \rightarrow strong and sparse signal.

In the presence of weak signals ($a = 0.02$), the Locfdr and Mixfdr show a great deal of variability, as shown in Fig 3 (first row). CDfdr maintains the smallest tail-specific MISE among all the methods, which again ensures its utility. For strong signals (second row of Fig 3), most of the methods have reasonable performance, except perhaps, ($a = 0.002, \pi_0 = 0.9$) case where the Locfdr poorly approximates the tail. The large number of outliers for Fdrtool are also not desirable.

Overall, it is encouraging that **CDfdr** *adapts to the underlying signal sparsity and strength in many cases*, which makes it very attractive and reliable for large-scale studies. Undoubtedly for the examples we have discussed in this paper it appears that CDfdr shows the most consistent and robust performance.

5 Conclusion

We have shown how the concepts and notations of comparison density and distribution (1) provide unification (and an alternative rationale) of the two cultures of multiple testing:

Frequentist BH procedure and Efron’s empirical Bayes local fdr idea; and (2) lead to a new class of efficient nonparametric signal detection algorithms. Our approach converts the mixed sample signal detection problem into a nonparametric comparison density function estimation problem. The author expects proposed CD-based nonparametric functional statistical approach will not only simplify the theory (which has an enormous literature) and practice but will also allow the possibility to include large-scale inference topics *as part of a conventional statistical inference curriculum* that can cover one-sample, mixed-sample and two-sample problems using *single concept, notion, and algorithm*. Our unified treatment by *linking* different large-scale multiple testing ideas (rather than presenting them as “unrelated tools”) could accelerate the learning for students.

We made a crucial observed that almost all large-scale inference problems produce comparison density with a typical stylized shape – “U” shaped density over the compact interval $[0, 1]$, which is not easily estimable using traditional nonparametric density estimation techniques. To address this, we introduced a *new genre* of density estimation methods based on the idea of Skew-G density decomposition. This allows for richer data-driven specification for tail-modeling via simple parametric models, which has an added advantage of being interpretable and easily implementable. We believe this technique can be used in many other modeling problems (outside multiple testing), such as heavy-tailed density estimation.

In future work, we would like to extend this technique to discrete mixed sample problems. Although we believe the theory presented in Mukhopadhyay and Parzen (2014) will guide us in that direction, it will require more detailed studies. Systematic theoretical investigation of the asymptotic properties of the proposed nonparametric skew-beta density estimator is an unexplored and open problem. Additionally, we plan to understand the effect of dependence on our proposed density estimator and how it influences the false discovery thresholds.

6 Supplementary Materials

Web Appendices, Tables and Figures referenced in Sections 1.1, 2.2, 3.2, 3.4 and 4 are available with this paper at the *Biometrics* website on Wiley Online Library. Online link: <http://bit.do/LSSD-BiometricsSupp>.

Acknowledgment

The author would like to thank Emanuel Parzen for several valuable comments and suggestions. The author also thank the Associate Editor and anonymous referees whose constructive comments have greatly helped to improve the quality and presentation of the paper. The author dedicates this paper to John Tukey, the pioneer of multiple comparison idea, on the occasion of his 100th birthday.

References

- Anderson, G. L. and de Figueiredo, R. J. P. (1980). An adaptive orthogonal series estimator for probability density functions. *Annals of Statistics* **8**, 347–376.
- Beaton, A. E. and Tukey, J. W. (1974). The fitting of power series, meaning polynomials, illustrated on band-spectroscopic data. *Technometrics* **16**, 147–185.
- Benjamini, Y. (2008). Comment: Microarrays, empirical bayes, and the two-groups model. *Statistical Science* **23**, 23–28.
- Benjamini, Y. (2010). Simultaneous and selective inference: current successes and future challenges. *Biometrical Journal* **52**, 708–721.
- Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J Roy Statist Soc Ser B.* **57**, 289–300.
- Donoho, D. and Jin, J. (2004). Higher criticism for detecting sparse heterogeneous mixtures. *The Annals of Statistics* **32**, 962–994.
- Efron, B. (2004). Large-scale simultaneous hypothesis testing: The choice of a null hypothesis. *Journal of the American Statistical Association* **99**, 96–104.
- Efron, B. (2007). Size, power and false discovery rates. *Annals of Statistics*. **35**, 1351–1377.
- Efron, B. (2010). *Large-scale inference: empirical Bayes methods for estimation, testing, and prediction*, volume 1. Cambridge; New York: Cambridge University Press.
- Efron, B., Storey, J., and Tibshirani, R. (2001). Microarrays, empirical Bayes methods, and false discovery rates. *Journal of the American Statistical Association* **96**, 1151–60.

- Geenens, G. (2014). Probit transformation for kernel density estimation on the unit interval. *Journal of the American Statistical Association* **109**, 346–358.
- Golub, T., Slonim, D., Tamayo, P., C. Huard, M. G., Mesirov, J., Coller, H., Loh, M., Downing, J., Caligiuri, M., Bloomfield, C., and Lander, E. (1999). Molecular classification of cancer: Class discovery and class prediction by gene expression. *Science* **286**, 531–537.
- Ledwina, T. (1994). Data driven version of neyman smooth test of fit. *Journal of the American Statistical Association* **89**, 1000–1005.
- Mukhopadhyay, S. and Parzen, E. (2014). LP approach to statistical modeling. *Preprint arXiv:1405.2601*.
- Muralidharan, O. (2010). An empirical Bayes mixture method for effect size and false discovery rate estimation. *Annals of Applied Statistics* **4**, 422–438.
- Parzen, E. (1983). Fun.stat quantile approach to two sample statistical data analysis. *Technical Report, Texas A&M University*.
- Parzen, E. (1999). Statistical methods mining, two sample data analysis, comparison distributions, and quantile limit theorems. In *Szyszkowicz, B., editor, Asymptotic Methods in Probability and Statistics*. Elsevier, Amsterdam. pages 611–617.
- Schweder, T. and Spjøtvoll, E. (1982). Plots of p-values to evaluate many tests simultaneously. *Biometrika* **69**, 493–502.
- Singh, D., Febbo, P. G., Ross, K., Jackson, D. G., Manola, J., Ladd, C., Tamayo, P., Renshaw, A. A., D’Amico, A. V., Richie, J. P., Lander, E. S., Loda, M., Kantoff, P. W., Golub, T. R., and Sellers, W. R. (2002). Gene expression correlates of clinical prostate cancer behavior. *Cancer Cell* **1**, 203–209.
- Storey, J. (2002). A direct approach to false discovery rates. *Journal of the Royal Statistical Society, Series B* **64**, 479–498.
- Strimmer, K. (2008). A unified approach to false discovery rate estimation. *BMC Bioinformatic* **9**, 1–14.
- Thas, O. (2010). *Comparing Distributions*. Springer, New York.

- Tukey, J. W. (1949). Comparing individual means in the analysis of variance. *Biometrics* **5**, 99–114.
- Tukey, J. W. (1989). Higher criticism for individual significances in several tables or parts of tables. *Internal working paper, Princeton University* .
- Wen, K. and Wu, X. (2014). An improved transformation-based kernel estimator of densities on the unit interval. *Journal of the American Statistical Association (in press)* **110**, 773–783.

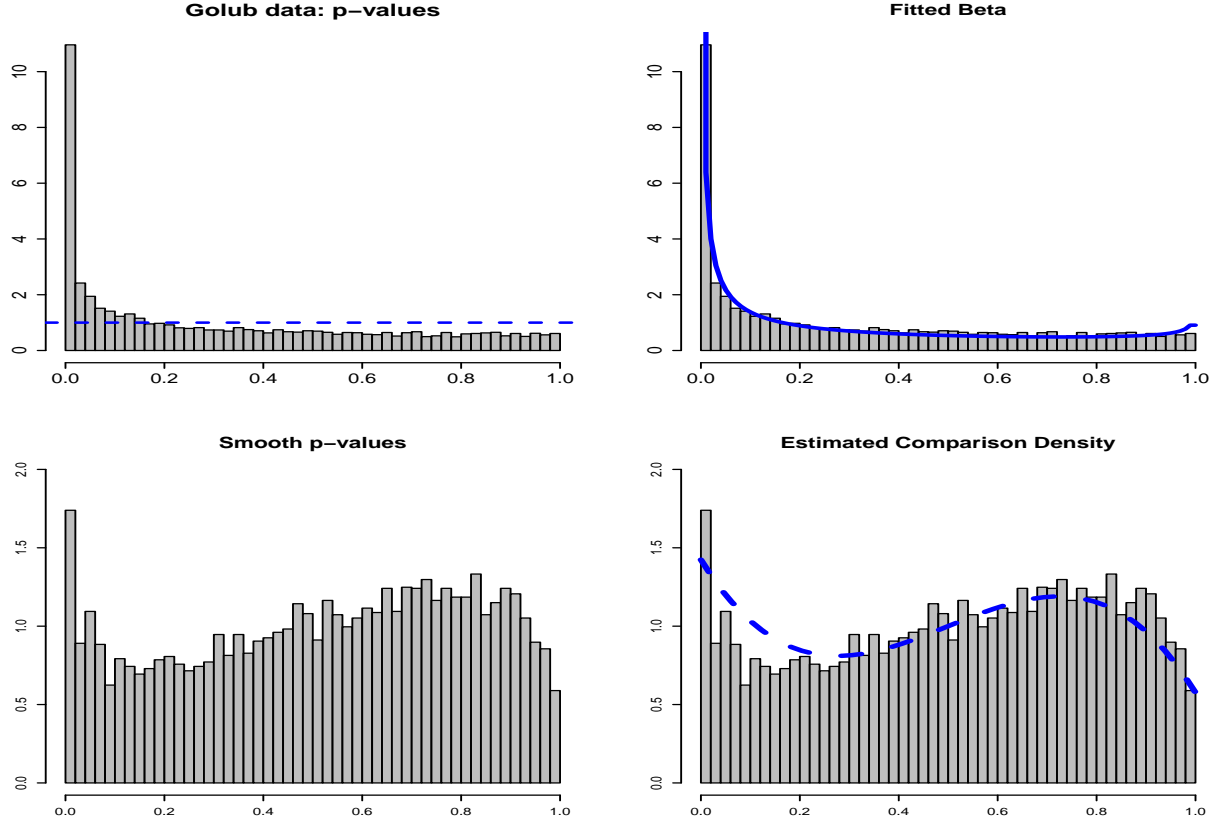


Figure 1: The concepts of ‘smooth’ p-value and the mechanism of skew-Beta density estimation for Golub gene expression data is illustrated; (A) [Top Left] Histogram of the 7129 p-values of using two sample t-test; (B) [Top Right] Fitted beta distribution $\text{Beta}(\hat{\alpha} = .32, \hat{\beta} = .75)$ to the p-values; (C) [Bottom Left] Distribution of the smooth p-values $v = F_B(u; \hat{\alpha} = .32, \hat{\beta} = .75)$; (D) [Bottom Right] LP-adaptive orthogonal (Nonparametric) density estimator of the smooth Golub p-values: $\hat{d}(v; F_B, F) = 1 - .16 \text{Leg}_3(v), 0 < v < 1$.

Comparing estimated fdr for various methods




Figure 2: (color online) Estimated local fdr is shown for seven different methods for Prostate cancer data.

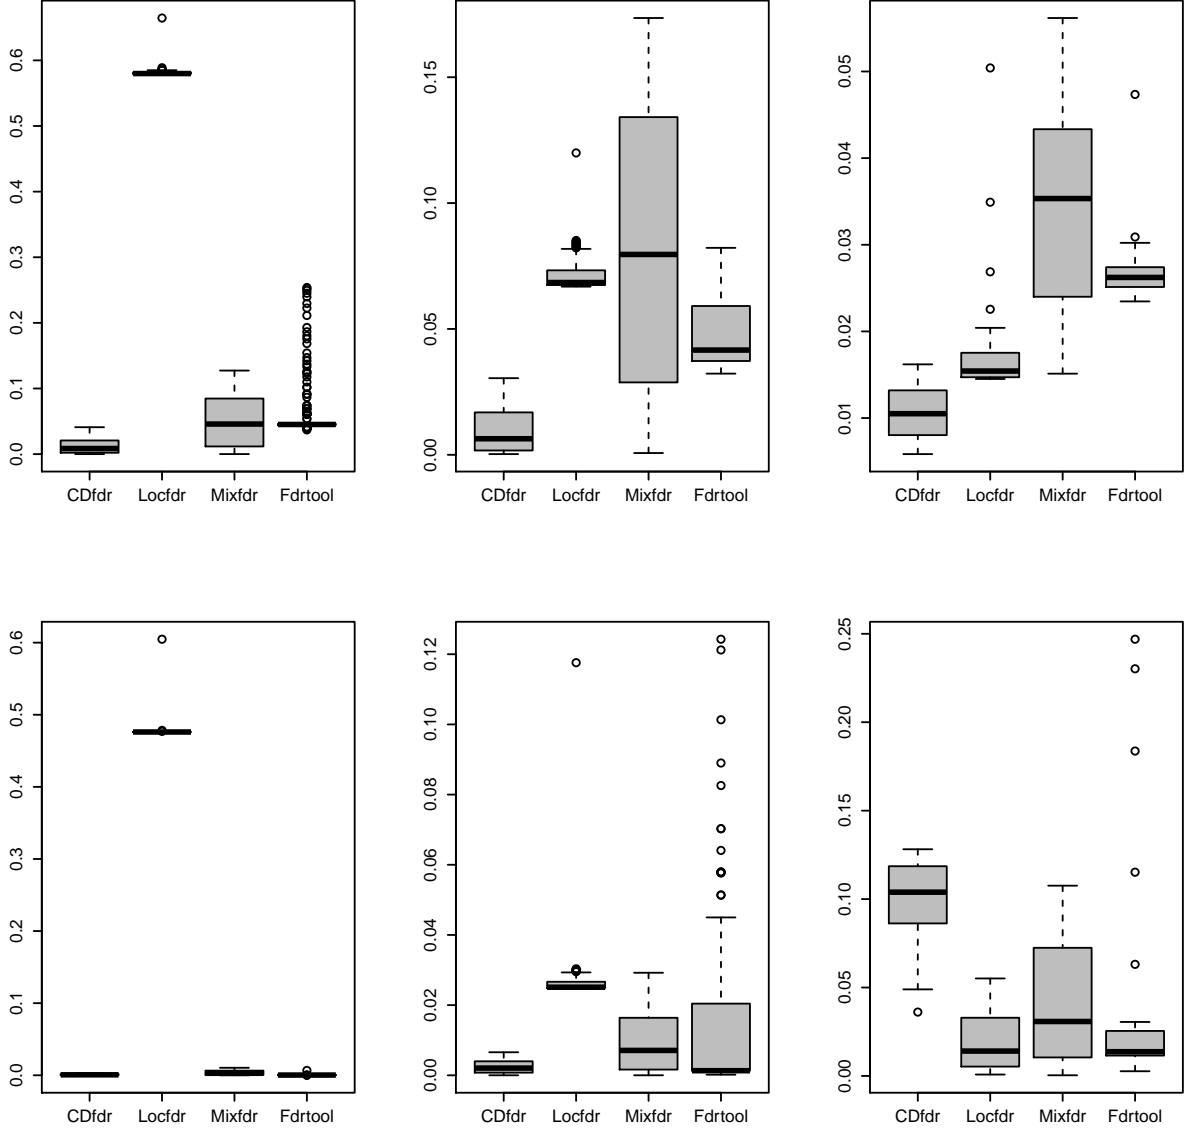


Figure 3: Compares the tail specific MISE for Uniform mixture model $\pi_0 \text{Uniform}[0, 1] + (1 - \pi_0) \text{Uniform}[0, a]$. The rows corresponds to $a = 0.02$ and 0.002 . For each row the columns (from left to right) denotes $\pi_0 = 0.9, 0.95, 0.99$.

Table 1: *Comparing three empirical null estimation methods on five large scale data sets. Readers are referred to Appendix B of Efron (2010) for more information on the data sets.*

Examples	Methods	$\hat{\mu}_0$	$\hat{\sigma}_0$
HIV	Robust biweight M-estimator	.121	.845
	MLE (Locfdr)	.115	.753
	Penalized mixture model (Mixfdr)	.131	.838
Gene-Tagging	Robust biweight M-estimator	.383	1.25
	MLE (Locfdr)	.310	1.26
	Penalized mixture model (Mixfdr)	.373	1.23
Leukemia	Robust biweight M-estimator	.017	1.88
	MLE (Locfdr)	.120	1.59
	Penalized mixture model (Mixfdr)	.017	1.86
Prostate	Robust biweight M-estimator	−0.001	1.092
	MLE (Locfdr)	−0.002	1.087
	Penalized mixture model (Mixfdr)	−0.002	1.068
NYC Police	Robust biweight M-estimator	.087	1.42
	MLE (Locfdr)	.121	1.39
	Penalized mixture model (Mixfdr)	.080	1.42